arXiv:1609.08377v2 [cond-mat.mtrl-sci] 6 Oct 2016

Extremely scalable algorithm for 10^8-atom quantum material simulation on the full system of the K computer

Takeo Hoshi and Hiroto Imachi
Department of Applied Mathematics and Physics
Tottori University
4-101 Koyama-Minami, Tottori, Japan
Email(Hoshi): hoshi@damp.tottori-u.ac.jp

Kiyoshi Kumahata, Masaaki Terai,
Kengo Miyamoto, Kazuo Minami and Fumiyoshi Shoji
Operations and Computer Technologies Division
RIKEN Advanced Institute for Computational Science
7-1-26 Minatojima-minami-machi, Chuo-ku, Kobe, Japan

***Abstract*—An extremely scalable linear-algebraic algorithm was developed for quantum material simulation (electronic state calculation) with 10^8 atoms or 100-nm-scale materials. The mathematical foundation is generalized shifted linear equations ($(zB-A)\boldsymbol{x}=\boldsymbol{b}$), instead of conventional generalized eigenvalue equations. The method has a highly parallelizable mathematical structure. The benchmark shows an extreme strong scaling and a qualified time-to-solution on the full system of the K computer. The method was demonstrated in a real material research for ultra-flexible (organic) devices, key devices of next-generation Internet-of-Things (IoT) products. The present paper shows that an innovative scalable algorithm for a real research can appear by the co-design among application, algorithm and architecture.**

I. Introduction

Large-scale quantum material simulation (electronic state calculation) is a major field of computational science and engineering. Calculations for one-hundred-million (10^8) atoms or 100-nano-meter(nm)-scale systems have a strong need for innovative industrial products but are far beyond the computational limit of the present standard methods. The present paper reports that a novel linear algebraic algorithm [1], [2], [3], [4], [5], [6] shows an extreme strong scaling and a qualified time-to-solution on the full system on the K computer with 10^8 atoms or 100-nm-scale systems. The algorithm was implemented in our code ELSES (=Extra-Large-Scale Electronic Structure calculation; http://www.elses.jp/). The method was demonstrated with condensed polymer systems that appears in an academic-industrial collaboration research for next-generation Internet-of-Things (IoT) devices.

The present paper is organized as follows; The background or the algorithm is presented in Sec. II or Sec. III, respectively. The benchmark and their analysis are given in Sec. IV. The application in real research is given in Sec. V. The conclusion is given in Sec. VI.

II. Background

A. Large-scale eigenvalue problem and its difficulty

A mathematical foundation of electronic state calculations is a generalized eigenvalue problem of

$$A\boldsymbol{y}_k = \lambda_k B\boldsymbol{y}_k. \tag{1}$$

The matrices A and B are Hamiltonian and the overlap matrices, respectively. These matrices are $M \times M$ Hermitian matrices and B is positive definite. In this paper, these matrices are real-symmetric. An eigenvalue (λ_k) or eigenvector ($\boldsymbol{y}_k$) represents the energy or quantum wavefunction $\phi_k(\boldsymbol{r})$ of one electron, respectively. In typical cases, the matrix size M is proportional to the number of atoms N ($M \propto N$).

Direct eigenvalue solvers consume $\mathrm{O}(M^3)$ operation costs and their practical limit is the matrix size of $M = 10^6$ for the current supercomputers. Recently, a million dimensional eigenvalue problem, the largest problem as far as we know, was solved by an optimally hybrid solver (EigenKernel; https://github.com/eigenkernel/) [5] with the two modern solvers of ELPA [7] and EigenExa [8]. The ELPA routine was used for the reducing procedure from the generalized eigenvalue problem into the standard one, while the EigenExa routine was used so as to solve the reduced standard problem. The elapsed time on the K computer is $T_{\mathrm{elaps}} = 9,939$ sec with $n_{\mathrm{node}} = 41,472$ nodes and $T_{\mathrm{elaps}} = 5,516$ sec on the full system (with $n_{\mathrm{node}} = 82,944$ nodes).

The large-scale problem of Eq. (1) has a potential difficulty, because an explicit orthogonalization procedure is required with $\mathrm{O}(N^3)$ operation costs, so as to satisfy the orthogonality relation of $\boldsymbol{y}_k^{\mathrm{T}} B \boldsymbol{y}_l = \delta_{kl}$. The above potential difficulty appears commonly among the large-scale electronic state calculations. A calculation code with $\mathrm{O}(N^3)$ operation costs is RSDFT [9], the winner of Gordon Bell Prize in 2011. The method is based on first principles and real-space mesh grid and was used with up to $N = 10^5$ atoms on the K computer. Although the above paper is a fascinating progress, the present target is far beyond the computational limit.

B. A novel concept for large-scale calculations

A novel concept for large-scale calculations was proposed by Walter Kohn, a winner of the Nobel Prize in Chemistry at 1998. His paper in 1996 shows that the above potential difficulty in electronic state calculation can be avoided, when the theory is not based on an eigenvalue problem and the formulation is free from the orthogonalization procedure [10]. The concept realizes 'order-N' methods [1], [11], [12], [13], [14], in which the computational cost is $\mathrm{O}(N)$ or proportional to the number of atoms N.

Here the concept [10] is briefly explained. The theory focuses on a physical quantity defined as

$$\langle X \rangle \equiv \sum_k f(\lambda_k)\, \boldsymbol{y}_k^{\mathrm{T}} X \boldsymbol{y}_k, \tag{2}$$

with a given sparse real-symmetric matrix X. Equation (2) is found in elementary textbooks of electronic state calculations. The function of $f(\lambda)$ is a weight function, called Fermi function, and is defined as

$$f(\lambda) \equiv \left\{ 1 + \exp(\frac{\lambda - \mu}{\tau}) \right\}^{-1}. \tag{3}$$

The weight function is a 'smoothed' step function with a smoothing parameter $\tau (> 0)$, because the Heaviside step function will appear in the limiting case of $\tau \to +0$ ($f(\lambda) = 1(\lambda < \mu)$ and $f(\lambda) = 0(\lambda > \mu)$). The smoothing parameter τ indicates the temperature of electrons. The parameter μ is the chemical potential and the value should be determined, so as to reproduce the number of electrons in the material. The case in $X = A$, for example, gives the electronic energy

$$\langle A \rangle \equiv \sum_k f(\lambda_k)\, \lambda_k. \tag{4}$$

A quantity in Eq.(2) is transformed into the trace form of

$$\langle X \rangle = \mathrm{Tr}[\rho X] = \sum_{i,j} \rho_{ji} X_{ij} \tag{5}$$

with the density matrix

$$\rho \equiv \sum_k f(\varepsilon_k)\, \boldsymbol{y}_k\, \boldsymbol{y}_k^{\mathrm{T}}. \tag{6}$$

The order-N property can appear, since the matrix X is sparse; A density matrix element ρ_{ji} is *not* required when $X_{ij} = 0$, because the element ρ_{ji} does not contribute to the physical quantity of Eq. (5), even if its value is nonzero ($\rho_{ji} \neq 0$). Consequently, the number of the required density matrix elements ρ_{ji} is $\mathrm{O}(N)$. The above fact is called 'quantum locality' or 'nearsightedness principle' [10].

The above formulation has a highly parallelizable mathematical structure and the original problem is decomposed mathematically into parallel subproblems. The trace in Eq. (5) can be decomposed as

$$\langle X \rangle = \mathrm{Tr}[\rho X] = \sum_j^M \boldsymbol{e}_j^{\mathrm{T}} \rho X \boldsymbol{e}_j, \tag{7}$$

with the j-th unit vector of $\boldsymbol{e}_j \equiv (0, 0, 0, ..., 1_j, 0, 0, ..., 0)^{\mathrm{T}}$. Here the quantity of $\boldsymbol{e}_j^{\mathrm{T}} \rho X \boldsymbol{e}_j$ is called 'projected physical quantity', because the quantity is defined by the projection onto the vector of $\boldsymbol{e}_j$. The essence of the parallelism is the fact that the projected physical quantity of $\boldsymbol{e}_j^{\mathrm{T}} \rho X \boldsymbol{e}_j$ is calculated almost independently among different indices of j.

An important application is quantum molecular dynamics simulation, in which an electron is treated as a quantum mechanical wave, while an atom (a nucleus) is treated as a classical particle in Newtonian equation of motion

$$M_I \frac{d^2 \boldsymbol{R}_I}{dt^2} = \boldsymbol{F}_I. \tag{8}$$

Here, M_I and $\boldsymbol{R}_I$ are the mass and the position of the I-th atom and $\boldsymbol{F}_I$ is the force on the I-th atom. Other variables, such as the electronic charge on each atom $\{q_I\}_I$, can be also calculated. The force and charge on each atom can be calculated in the trace form of Eq. (5).

C. Physical origin of the matrices

The matrices of A and B are sparse and their physical origin is found in Ref. [1] and reference therein. In short, the calculations in the present paper are formulated by a first-principle-based modeled (transferable tight-binding) theory. An electronic wavefunction $\phi_k(\boldsymbol{r})$ is expressed by an eigenvector of $\boldsymbol{y}_k \equiv (y_{1k}, y_{2k}, ..., y_{Mk})^{\mathrm{T}}$, as $\phi_k(\boldsymbol{r}) = \sum_j y_{jk} \chi_j(\boldsymbol{r})$ with the given basis functions of $\{\chi_j(\boldsymbol{r})\}_j$ called atomic orbitals. A basis function is localized in real space and its localization center is the position of one atom. A matrix element of A_{ij} or B_{ij} represents the quantum (wave) interaction of electrons on the i-th and j-th bases. The basis index, i or j, is the composite indices of the atom index I or J that distinguishes the localization center and another index, α or β, called orbital index that distinguishes the shape of the function ($i \Leftrightarrow (I, \alpha), j \Leftrightarrow (J, \beta)$). An element of the matrices A and B can be expressed by the four indices as $A_{I\alpha;J\beta}$ and $B_{I\alpha;J\beta}$, respectively. The matrices are sparse, because the elements decays quickly ($|A_{I\alpha;J\beta}|, |B_{I\alpha;J\beta}| \to 0$) as the function of the distance between the I-th and J-th atoms (r_{IJ}). In the present simulation, a cutoff distance r_{cut} was introduced so that a matrix element, $A_{I\alpha;J\beta}$ or $B_{I\alpha;J\beta}$, is ignored in the cases of $r_{IJ} > r_{\mathrm{cut}}$. Among the present benchmarks, the cutoff distance r_{cut} is set to be $r_{\mathrm{cut}} = 5$ au (≈ 0.2646nm) for diamond crystal and $r_{\mathrm{cut}} = 10$ au (≈ 0.5292nm) for condensed polymers. A longer cutoff distance is used for condensed polymers, so as to include the interaction between polymers.

The number of orbitals on one atom can be different among atom species. The simulated materials in the present paper consists in hydrogen (H) and carbon (C) atoms. One (s-type) orbital is prepared at each hydrogen (H) atom, and four (s-, p_x-, p_y-, p_z-types) atomic orbitals at each carbon (C) atom. A material with N_{H} hydrogen atoms and N_{C} carbon atoms gives the matrices of A and B with the size of $M = N_{\mathrm{H}} + 4N_{\mathrm{C}}$.

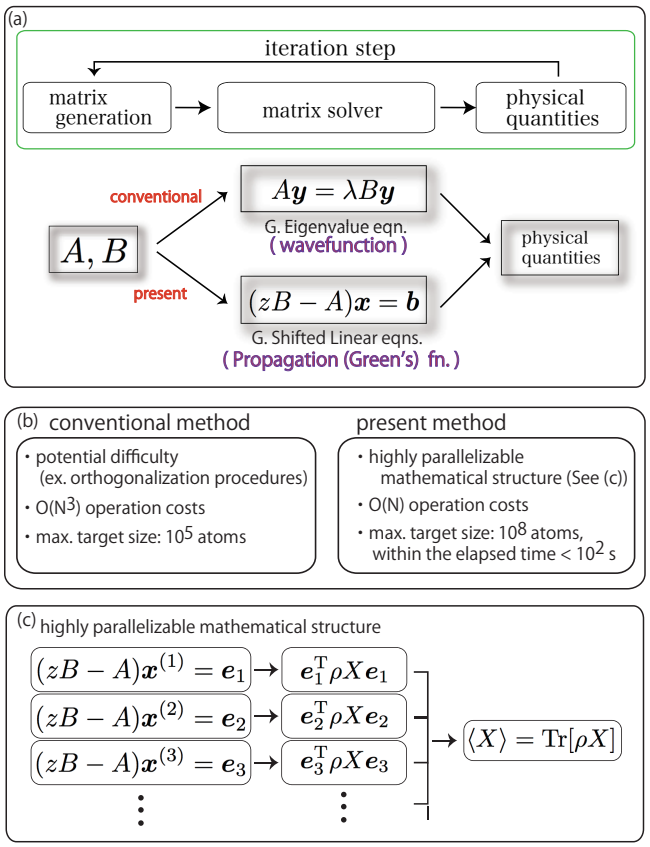

Fig. 1: (a) The ground design of scalable algorithm. (b) Comparison between the conventional and present method. (c) Illustration of the highly parallelizable mathematical structure.

III. Algorithm

A. Ground design

The ground design of the present scalable algorithm [1] is shown in Fig. 1(a). The comparison between the conventional and present methods is summarized in Fig. 1(b). The method is based not on the eigenvalue problem of Eq. (1) but on the set of linear equations in the form of

$$(zB - A)\boldsymbol{x} = \boldsymbol{b}. \tag{9}$$

Here, z is a complex energy value. The vector $\boldsymbol{b}$ is an input and the vector $\boldsymbol{x}$ is the solution vector. A set of linear equations in the form of Eq. (9) with different energy values ($z = z_1, z_2, ...$) is called generalized shifted linear equations. The case in $B = I$ is called shifted linear equations. The use of Eq. (9) results in the Green's (propagation) function formalism, since the solution $\boldsymbol{x}$ of Eq. (9) is written formally as

$$\boldsymbol{x} = G\boldsymbol{b} \tag{10}$$

with the Green's function $G \equiv (zB - A)^{-1}$. The Green's function and the eigenvectors holds the relationship of

$$G(z) = \sum_k^M \frac{\boldsymbol{y}_k\, \boldsymbol{y}_k^{\mathrm{T}}}{z - \lambda_k}. \tag{11}$$

The density matrix is also given by the Green's function as

$$\rho = \frac{-1}{\pi}\int_{-\infty}^{\infty} f(\varepsilon)\,\mathrm{Im}[G(\varepsilon + i0)]\,d\varepsilon. \tag{12}$$

The present method has a highly parallelizable mathematical structure, as illustrated in Fig. 1(c), since the projected physical quantity of $\boldsymbol{e}_j^{\mathrm{T}}\rho X\boldsymbol{e}_j$ in Eq. (7) is obtained from the generalized shifted linear equations of

$$(zB - A)\boldsymbol{x}^{(j)} = \boldsymbol{e}_j. \tag{13}$$

B. Krylov subspace solver

The generalized shifted linear equations of Eq. (13) are solved on an iterative Krylov-subspace solver. A Krylov subspace is defined as the linear space of

$$K_\nu(Q; \boldsymbol{b}) \equiv \mathrm{span}[\boldsymbol{b}, Q\boldsymbol{b}, Q^2\boldsymbol{b}, ..., Q^{\nu-1}\boldsymbol{b}], \tag{14}$$

with a given vector $\boldsymbol{b}$ and a given square matrix Q. An example is Conjugate Gradient method and the subspace dimension of ν is the number of iterations. Krylov-subspace methods with (generalized) shifted linear equations have been investigated in particular from 2000's, partially because the strategy is suitable to parallelism. Since the solver algorithms are mathematical, they are applicable to many scientific areas, such as, QCD [15], large-scale electronic state calculation [1], [16], [17], [18], [19], quantum many-body electron problem [20], nuclear shell model problem [21], first-principle electronic excitation problem [22], and first-principle transport calculation [23]. In the present paper, the multiple Arnoldi solver [1] is used, in which Eq. (13) is solved within the direct sum of the two Krylov subspaces of

$$\mathcal{L}_\nu(A, B; \boldsymbol{e}_j) \equiv K_{\nu/2}(A; \boldsymbol{e}_j) \oplus K_{\nu/2}(A; B^{-1}\boldsymbol{e}_j) \tag{15}$$

with an even number of ν. The number ν is typically, $\nu = 30 - 300$ and the calculations in the present paper was carried out with $\nu = 30$ as in the previous one [1]. The second term in the right hand side of Eq. (15) appears so as to satisfy several conservation laws [1]. A reduced (small) $\nu \times \nu$ eigenvalue equation is solved and the solution vector is given by

$$\boldsymbol{x}^{(j)} := G^{(j)}(z)\boldsymbol{e}_j \tag{16}$$

with

$$G^{(j)}(z) \equiv \sum_m^\nu \frac{\boldsymbol{v}_m^{(j)}\, \boldsymbol{v}_m^{(j)\mathrm{T}}}{z - \varepsilon_m^{(j)}}. \tag{17}$$

Here $\varepsilon_m^{(j)}$ and $\boldsymbol{v}_m^{(j)}$ is an eigenvalue and eigenvector of the reduced equation ($m = 1, 2, ...\nu$). When the Green's function of G in Eq. (12) is replaced by $G^{(j)}(z)$ in Eq. (17), the projected physical quantity with the index of j is given by

$$\begin{aligned} \boldsymbol{e}_j^{\mathrm{T}}\rho X\boldsymbol{e}_j &:= \frac{-1}{\pi}\int_{-\infty}^{\infty} f(\varepsilon)\,\mathrm{Im}[\boldsymbol{e}_j^{\mathrm{T}} G^{(j)}(\varepsilon + i0) X\boldsymbol{e}_j]\,d\varepsilon \\ &= \sum_m^\nu f(\varepsilon_m^{(j)})\boldsymbol{e}_j^{\mathrm{T}}\boldsymbol{v}_m^{(j)}\, \boldsymbol{v}_m^{(j)\mathrm{T}} X\boldsymbol{e}_j. \end{aligned} \tag{18}$$

An advantage of the method is that the energy integration is carried out analytically as in Eq.(18). Equation (18) will be

exact, if the subspace dimension of ν increases to the original matrix dimension ($\nu = M$). As an additional technique in large-scale calculations, the real-space projection technique [1] was also used. The radius of the spherical region is determined with an input integer parameter κ, so that the region contains κ atoms or more. The same technique is used also for the overlap matrix B. The value of κ is set to $\kappa = 100$ in the present paper as in the previous one [1]. As results, numerical problems in the form of Eq. (13) are solved with the matrix size of, typically, $M' = 200 - 400$ in the present paper.

C. Implementation

The code is written in Fortran 90 with the MPI/OpenMP hybrid parallelism. According to the parallel scheme in Fig. 1(c), the projected physical quantity of $\boldsymbol{e}_j^{\mathrm{T}} \rho X \boldsymbol{e}_j$ is calculated as single-thread or single-core calculations. As explained in Sec. II-C, the basis index j is a composite suffix of the atom index J and the orbital index β ($j \Leftrightarrow (J, \beta)$) In the code, the loop for the basis index j is implemented as the double loop that consists of the outer loop for the atom index J and the inner loop for the orbital index β. Since the outer loop is parallelized both in MPI and OpenMP parallelism, a meaningful parallel computation is possible, when the number of atoms is larger than that of cores ($N > n_{\mathrm{core}}$). Several matrix elements of A, B are generated redundantly among nodes, so as to save inter-node communications. The pure MPI parallelism is possible but consumes larger memory costs.

The communication among nodes is required, *only when* a summation is performed in the trace form of Eq. (7), as shown in Fig. 1(c) [1]. The summation is carried out hierarchically; First, the summation is carried out on each node by OpenMP directives and then the summation is carried out between nodes by `MPI_Allreduce()`.

IV. Benchmark

A. Purpose and condition

The benchmarks were carried out so as to show an extreme strong scaling and a time-to-solution qualified for a real research. Our target value of the qualified time-to-solution is $T_{\mathrm{elaps}} = 10^2$s for the elapsed time per step in a quantum molecular dynamics simulation, because a dynamical simulation of $n_{\mathrm{step}} = 10^3$ steps can be executed within one day ($T_{\mathrm{elaps}} n_{\mathrm{step}} = 10^5$s $\approx$ one day). The calculations were carried out on the K computer which consists of 82,944 compute nodes and achieved the peak performance of 11.28PFLOPS. Each CPU has eight cores and the interconnection between nodes is named 'Tofu' which constructs physical six-dimensional mesh/torus network topology. We used the `MPI_Allreduce()` optimized on the K computer [24].

The calculations were executed in double precision with the MPI/OpenMP hybrid parallelism. The number of the MPI processes is set to that of the compute nodes and the number of the OpenMP threads is set to be eight, the number of cores per compute node. The jobs were executed by specifying the three-dimensional node geometry on the K computer ($n_{\mathrm{node}} \equiv n_{\mathrm{node}}^{(x)} \times n_{\mathrm{node}}^{(y)} \times n_{\mathrm{node}}^{(z)}$) for optimal performance or minimum hop count. The number of used nodes (node geometry) is listed below; $n_{\mathrm{node}} = 2,592 (= 12 \times 12 \times 18)$, $5,184 (= 12 \times 18 \times 24)$, $10,368 (= 18 \times 24 \times 24)$, $20,736 (= 24 \times 27 \times 32)$, $41,472 (= 27 \times 32 \times 48)$, and $82,944 (= 32 \times 48 \times 54$, the full system).

The benchmark were carried out for disordered materials that appears in real research of ultra-flexible devices. Condensed polymer systems of poly-(phenylene-ethynylene) (PPE) were simulated. The three systems are called 'P100', 'P10' and 'P1' and contain N=101,606,400 ($\approx 10^8$ or 100M), N=10,137,600 ($\approx$ 10M) and N=1,228,800 ($\approx$1M) atoms, respectively. The periodic boundary condition is imposed. The size of the periodic simulation box is 134 nm $\times$ 134 nm $\times$ 209 nm for the 'P100' system. The simulations were carried out also for the ideal diamond solid called 'D100' that contains $N = 106,168,320$ ($\approx 10^8$ or 100M) atoms in the ideal periodicity, so as to discuss the influence of the presence or absence of structural disorder.

Technical details are explained. The initial atomic structures for the polymer systems were generated in classical molecular dynamics simulations by GROMACS (http://www.gromacs.org/). Classical simulations work faster but do not treat electronic (quantum) waves responsible for the device property. The recorded elapsed time was one for a 'snapshot' simulation, an electronic state calculation of the given atomic structures, which dominates the elapsed time in molecular dynamics simulations. A molecular dynamics simulation can not be carried out with $N = 10^8$ atoms, because the required memory size exceeds the limit of the K computer (16GB per node). The present snapshot calculation with $N = 10^8$ atoms consumes 9 GB per node and a molecular dynamics simulation requires a larger memory size, so as to store additional variables like velocity, force and so on. The benchmark of molecular dynamics simulation with $N = 10^7$ atoms will be discussed in the last paragraph of this section.

B. Result

The measured elapsed time is summarized in Table I. Here the parallel efficiency ratio α is defined by

$$\alpha \equiv (T_{\mathrm{elaps}}(n_0)/T_{\mathrm{elaps}}(n_{\mathrm{node}}))/(n_{\mathrm{node}}/n_0) \qquad (19)$$

with $n_0 \equiv 2,592$. For example, the parallel efficiency ratio α with 10^8 atoms and the maximum number of nodes ($n_{\mathrm{node}} = 82,944$) is $\alpha = 0.92$ for 'D100' and $\alpha = 0.75$ for 'P100'.

TABLE I: The measured elapsed times T_{elaps} (sec) for ideal diamond solid with 10^8 atoms ('D100') and condensed polymer systems with 10^8 atoms ('P100'), with 10^7 atoms ('P10') and with 10^6 atoms ('P1'). The ideal or measured speed-up ratio is shown inside the parenthesis.

n_{node}	D100	P100	P10	P1
2,592 (1)	1001.4 (1)	741.1 (1)	81.4 (1)	10.3 (1)
5,184 (2)	502.2 (1.99)	378.5 (1.96)	43.7 (1.86)	5.95 (1.73)
10,368 (4)	252.6 (3.96)	195.2 (3.80)	24.3 (3.35)	3.28 (3.14)
20,736 (8)	127.9 (7.83)	103.0 (7.19)	11.4 (7.14)	1.96 (5.26)
41,472 (16)	65.6 (15.3)	57.1 (13.0)	6.32 (12.9)	1.25 (8.23)
82,944 (32)	34.1 (29.4)	30.9 (24.0)	3.60 (22.6)	0.84 (12.2)

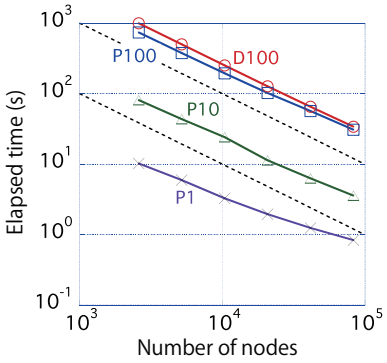

Fig. 2: Strong scaling benchmarks for ideal diamond solid with 10^8 atoms ('D100') and the condensed polymer systems with 10^8 atoms ('P100'), with 10^7 atoms ('P10') and with 10^6 atoms ('P1'). Dashed lines are drawn for ideal scaling.

Figure 2 shows the strong scaling property by plotting the data of Table I. In all the cases, the elapsed time $T_{\rm elaps}$ decreases monotonically as the function of the number of used nodes. The order-N property ($T_{\rm elaps} \propto N$) is also found. For example, the time of 'P100' is ten times larger that of 'P10' with $n_{\rm node} = 2,592$. As a rough estimation from Fig. 2, the target time-to-solution of $T_{\rm elaps} \approx 10^2$s is fulfilled by $n_{\rm node} \approx 2 \times 10^4$ and 2×10^3 for the condensed polymer systems with $N = 10^8$ and 10^7 atoms, respectively. The two cases conclude commonly that the qualified time-to-solution is fulfilled, when the number of atoms per node is approximately 5×10^2 ($N/n_{\rm node} \approx 5 \times 10^2$). The above statement can be interpreted as the weak-scaling property.

C. Analysis and discussion

Table II shows the measured communication and barrier times. In the simulations, we recorded not only the total elapsed time $T_{\rm elaps}$, but also the accumulated MPI communication time $T_{\rm comm}$ and the accumulated barrier time $T_{\rm barr}$ on all nodes. The barrier time includes the time to wait for other processors. The communication time $T_{\rm comm}$ is consumed by inter-node data communications, while the barrier time $T_{\rm barr}$ appears from a load imbalance among nodes.

Figure 3 plots the data in Tables I and II. Two points are discussed; (i) The communication time is not serious among all the cases. (ii) When the cases of 'D100' and 'P100' are compared, the ratio of the barrier time is much larger than

TABLE II: Communication time $T_{\rm comm}$ and barrier time $T_{\rm barr}$ of the elapsed time $T_{\rm elaps}$. See the caption of Table I for notations. The values are listed as $T_{\rm comm}$ / $T_{\rm barr}$ (sec).

$n_{\rm node}$	D100	P100	P10	P1
2,592	1.04 / 7.16	1.60 / 28.14	0.382 / 6.69	0.0617 / 2.68
5,184	1.04 / 3.34	1.60 / 20.75	0.378 / 4.86	0.0689 / 1.85
10,368	1.05 / 2.22	1.61 / 14.97	0.384 / 4.34	0.0734 / 1.13
20,736	1.05 / 1.34	1.61 / 9.05	0.218 / 3.31	0.0712 / 0.674
41,472	1.06 / 1.03	1.64 / 6.87	0.215 / 2.18	0.0727 / 0.409
82,944	1.06 / 0.485	1.65 / 5.96	0.218 / 1.28	0.0613 / 0.227

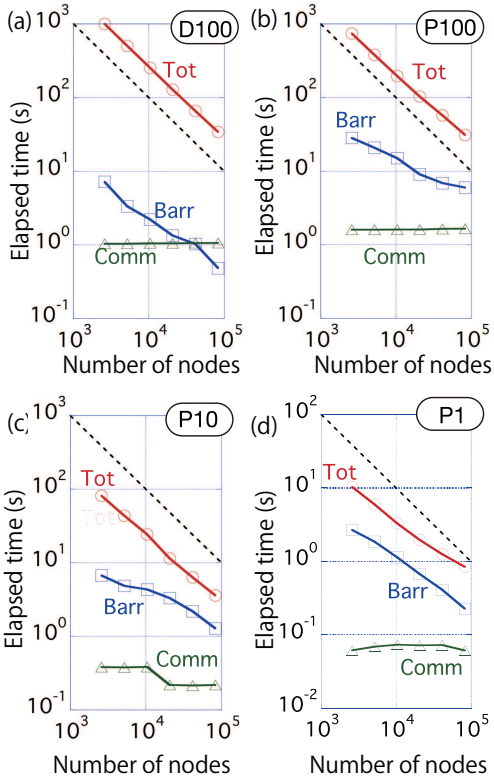

Fig. 3: Details of the elapsed time. The total elapsed time $T_{\rm elaps}$ (Tot), the communication time $T_{\rm comm}$ (Comm) and the barrier time $T_{\rm barr}$ (Barr) are plotted. See Fig. 2 for notations.

in 'P100'. In the full system calculation ($n_{\rm node} = 82,944$), for example, the ratio is $T_{\rm barr}/T_{\rm elaps} \approx 0.19$ in 'P100' and is ≈ 0.014 in 'D100'. We should recall that the 'D100' case is an ideal system without structural disorder and all the subproblems in Fig. 1(c) are equivalent. On the other hand, the load imbalance among nodes appears in 'P100', because of the structural disorder. The same conclusion holds on the 'P10' and 'P1' cases. A method for better load balance is a future (not urgent) issue of the present code.

To end up this section, two comments are addressed; (I) The further tuning should be focused mainly on single-core calculations, since the most routines are executed as single-core calculations as in Fig.1(c). The profiler reported that the performance is 2.3 % of the peak for the 'P100' case with $n_{\rm node} = 82,944$ in Table I. The severest limitation in the present calculations is the memory size of the K computer (16GB per node) and the present code was written in the memory-saving style, in which the memory cost should be minimized and the time cost is sometimes sacrificed. Since the situation can differ among materials and/or architectures, a possible way is to add another workflow in the time-saving style. The routines can be classified into those for the generation of matrix elements and and for the Krylov subspace solver as in Fig.1(a). The matrix-vector multiplication gives a

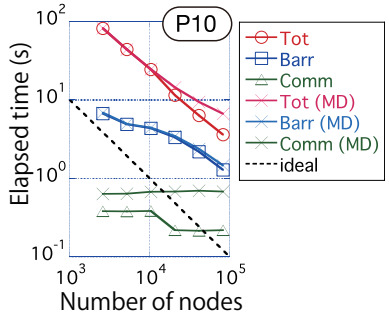

Fig. 4: Details of the elapsed time for the MD simulation in the 'P10' case. The total elapsed time $T_{\rm elaps}^{\rm (MD)}$ (Tot(MD)), the barrier time $T_{\rm barr}^{\rm (MD)}$ (Barr(MD), and the communication time $T_{\rm comm}^{\rm (MD)}$ (Comm(MD)) are plotted per MD step in the same manner of Fig. 3(c). The data for the electronic state calculation (Tot, Barr, Comm) are also plotted for comparison.

large fraction of the total elapsed time, as usual in a Krylov-subspace solver, and a typical fraction is 21 % among the present condensed polymer systems. The result suggests that the matrix generation part gives a larger fraction. (II) Fig. 4 shows the benchmark of molecular dynamics simulation for the 'P10' case, the possible maximum size (See the first paragraph of the present section), in the same manner of Fig. 3(c). For example, the elapsed time per molecular dynamics time step is $T_{\rm elaps}^{\rm (MD)}$ = 81.8 sec or 6.62 sec in $n_{\rm node}$ =2,592 or 82,944, respectively. For comparison, Fig. 4 also shows the data in Fig. 3(c), the data with the electronic structure calculation part. The elapsed time is much smaller than the target time-to-solution (10^2s) and the method is qualified well for a real research. We found, however, that non-negligible time costs appear in the total elapsed time ($T_{\rm elaps}^{\rm (MD)}$) among the cases with $n_{\rm node} > 2 \times 10^4$, because of the additional routine for MD simulation. Now we are tuning the code for faster MD simulations.

V. Application in Real Material Research

This section is devoted to the application study of the present method to a condensed polymer system, so as to show how a real research works well with $N = 10^8$ atoms by distributed computing. As an application study with a smaller system, a molecular dynamics simulation with $N = 10^5$ atoms was carried out with 10^4 cores and the elapsed time is 10 hours for 5,000 iteration steps [6]. Such a dynamical simulation is impractical with $N = 10^8$ atoms at the present day and this section indicates a part of the possible future research.

Here, the condensed organic polymer system of 'P100' was used. The research is motivated by an academic-industrial collaboration with Sumitomo Chemical Co., Ltd. [1], [3], [4]. Organic material gives the foundation of ultra-flexible (wearable) devices, key devices of next-generation IoT products, such as display, sensor and battery. A recent example is 'e-skin' [25]. The material is ultra-flexible (soft) and disordered in structure and the thickness of devices is typically 10^3 nm and 100-nm-scale simulations are crucial.

An important HPC issue is that the distributed data structure should be preserved throughout the whole research; Since the simulation data is huge and distributed, we cannot gather them into one node. Here we will show that the post-simulation data analysis works well for distributed data.

Figures 5(a)(b) show partial regions of the system and one can observe that the structure is fairly disordered. The molecular structure for a polymer unit is shown in Fig. 5(c). In general, electronic wavefunctions are localized in a disordered structure. The electrical current can propagate among polymers that are 'connected' locally by characteristic (π-type) electronic waves. We should investigate, therefore, the network of connected polymers.

A. Network analysis of electronic wavefunctions

A large-scale post-simulation data analysis was carried out so as to characterize the propagation of electronic wave in the disordered structure. A speculated propagation mechanism is shown schematically in Fig. 5(d). Three polymers are drawn and atoms are depicted as filled circles. The figures include a small local network that consists of two polymers connected by a dashed line. Electron can propagate along connected polymers. Since the network structure is dynamically changed, as schematically shown in Fig. 5(d), electron can propagate through the whole material.

The purpose of the analysis is to detect local polymer networks in which electronic wave can propagate. The analysis was carried out with the Green's function G obtained by the parallel order-N simulation, as follows; Stage I: The present parallel simulation gives a 'connectivity' matrix of C_{IJ}

$$C_{IJ} \equiv \sum_{\alpha}\sum_{\beta} \rho_{I\alpha;J\beta} H_{J\beta;I\alpha} \qquad (20)$$

where I, J are the atom indices. The connectivity matrix is called integrated crystal orbital Hamiltonian population (ICOHP) among physics papers [26], [2]. The quantity is a partial sum of the electronic energy $\langle H \rangle$ in Eq. (4) ($\langle H \rangle = \sum_{IJ} C_{IJ}$). Since the matrix elements are calculated always during the parallel simulation, the elements can be obtained independently among nodes, without any additional operation or communication cost. A matrix element C_{IJ} has a physical meaning of a local bonding energy between the I-th and J-th atoms; If the value of $|C_{IJ}|$ is significantly large, the two atoms are 'connected' by electronic wave. Stage II: Since every atom belongs to one of polymers, the connectivity matrix for polymers is defined by

$$C_{PQ}^{\rm (poly)} \equiv \sum_{I}^{\in P}\sum_{J}^{\in Q} C_{IJ}, \qquad (21)$$

where the summation of $\sum_I^{\in P}$, for example, means the summation among the atoms that belong to the P-th polymer. If an element $C_{PQ}^{\rm (poly)}$ shows a meaningful non-zero value, the P-th and Q-th polymers are connected by electronic wave. The

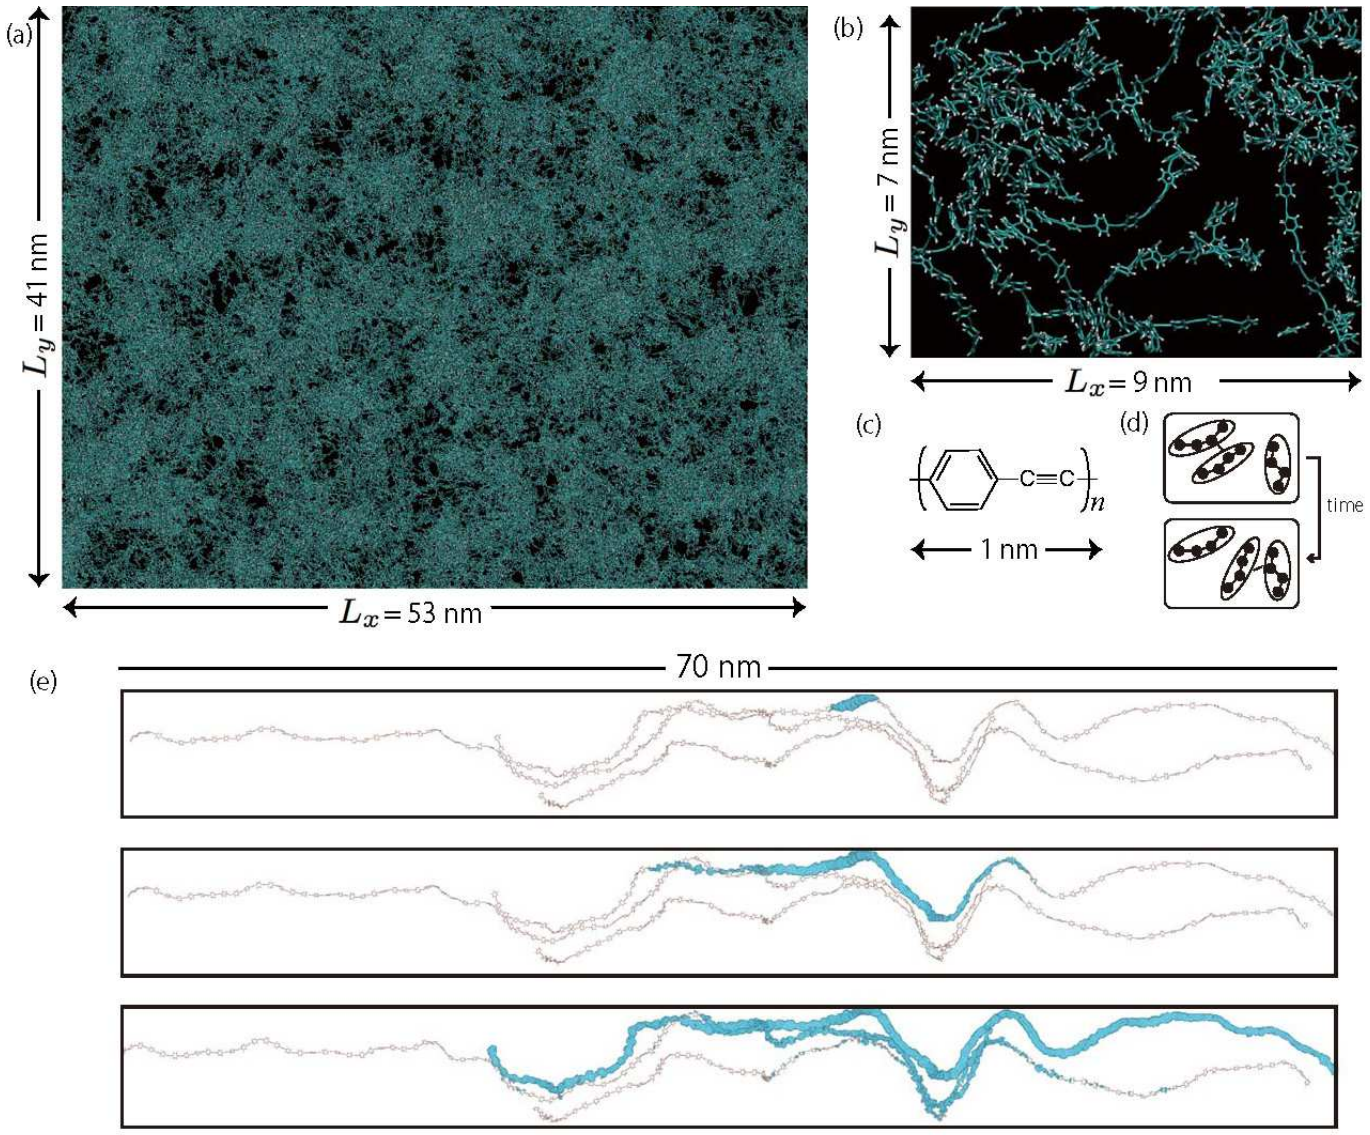

Fig. 5: A real material research for a condensed polymer system (PPE) with 10^8-atoms. (a) (b) Visualization of partial regions. The whole system has the periodic cell lengths of (265nm, 206nm, 239nm). (c) The unit structure of the polymer (PPE). (d) Schematic figures of a dynamically changed local network of connected polymers. (e) Quantum wave dynamics simulation of electrons with a local network of three connected polymers. The charge density of $q(\boldsymbol{r},t) \equiv |\Psi(\boldsymbol{r},t)|^2$ is drawn in the upper ($t = 0$), middle($t = 50$fs) and lower ($t = 948$fs) panels.

matrix $C^{(\mathrm{poly})}$ is sparse. The dimension of $C^{(\mathrm{poly})}$ is equal to the number of polymers $N^{(\mathrm{poly})} = 83,349$ and is much smaller than that of C ($N = 10^8$). Stage III: As a coarse grained analysis, the eigenvalue equation of $C^{(\mathrm{poly})}\boldsymbol{z} = \lambda \boldsymbol{z}$ in the matrix dimension of $N^{(\mathrm{poly})}$ was solved by the parallel eigenvalue solver [5]. As results, several eigenvectors $\boldsymbol{z}$ have several non-zero elements, which means the presence of small local networks with several connected polymers.

The network analysis reveals that the condensed polymer system has small networks that consist of several polymers, as illustrated in Fig. 5(d).

B. Quantum wave dynamics simulation for device property

Quantum wave (wavepacket) dynamics simulation [6] was carried out for device property, so as to confirm that the above network analysis is fruitful or that an electronic wave can propagate in the small polymer networks detected in the above analysis. In the wave dynamics simulation, an electronic wave $\Psi(\boldsymbol{r},t)$, a complex scalar vector, propagates dynamically under a Schrödinger-type equation of $i\partial_t \Psi = H\Psi$ with an effective Hamiltonian (matrix) H. See Ref. [6] and the references therein for details. The atom positions also change dynamically. Since the norm $q(\boldsymbol{r},t) \equiv |\Psi(\boldsymbol{r},t)|^2$ is the charge distribution, its dynamics gives the charge propagation or the

(non-stationary) electrical current. Figure 5(e) shows a typical dynamical simulation for approximately 1 ps. The simulation consumes six hours with 128 nodes. The result shows that the electronic wave propagates within a polymer first, and later propagates into other polymers, as expected. The method for large-scale wave dynamics simulation is under way.

VI. Conclusion

A novel linear algebraic algorithm realizes 10^8 atom or 100-nm-scale quantum material simulations with an extreme scalability and a qualified time-to-solution on the full system of the K computer. The mathematical foundation is generalized shifted linear equations, instead of conventional generalized eigenvalue equations and has a highly parallelizable mathematical structure. The method was demonstrated in a real material research for next-generation IoT products. The present paper shows that an innovative scalable algorithm for a real research can appear by the co-design among application, algorithm and architecture.

Acknowledgment

Several atomic structure data were provided from Masaya Ishida (Sumitomo Chemical Co., Ltd.). Part of the results is obtained by using the K computer at the RIKEN Advanced Institute for Computational Science (Proposal number hp150144, hp150281, hp160066, hp160222). This research is partially supported by Japan Science and Technology Agency, Core Research for Evolutional Science and Technology (JST-CREST) in the research area of 'Development of system software technologies for post-peta scale high performance computing'. This research is partially supported also by Grant-in-Aid for Scientific Research (KAKENHI Nos. 26400318 and 26286087, 16KT0016) from the Ministry of Education, Culture, Sports, Science and Technology (MEXT) of Japan. A part of the research is based on the collaboration with Priority Issue (Creation of new functional devices and high-performance materials to support next-generation industries) to be tackled by using Post 'K' Computer, MEXT, Japan.

References

[1] T. Hoshi, S. Yamamoto, T. Fujiwara, T. Sogabe, and S.-L. Zhang, "An order-N electronic structure theory with generalized eigenvalue equations and its application to a ten-million-atom system," J. Phys. Condens. Matter, 21, 165502, 2012.

[2] T. Hoshi, Y. Akiyama, T. Tanaka, and T. Ohno, "Ten-million-atom electronic structure calculations on the K computer with a massively parallel order-N theory," J. Phys. Soc. Jpn. 82, 023710, 2013.

[3] T. Hoshi, K. Yamazaki, and Y. Akiyama, "Novel linear algebraic theory and one-hundred-million-atom electronic structure calculation on the K computer," JPS Conf. Proc. 1, 016004, 2014.

[4] T. Hoshi, T. Sogabe, T. Miyata, D. Lee, S.-L. Zhang, H. Imachi, Y. Kawai, Y. Akiyama, K. Yamazaki, and S. Yokoyama, "Novel linear algebraic theory and one-hundred-million-atom quantum material simulations on the K computer," PoS(IWCSE2013) 065, 2014.

[5] H. Imachi and T. Hoshi, "Hybrid numerical solvers for massively parallel eigenvalue computation and their benchmark with electronic structure calculation," J. Inf. Process. 24, pp. 164–172, 2016.

[6] H. Imachi, S. Yokoyama, T. Kaji, Y. Abe, T. Tada, and T. Hoshi, "One-hundred-nm-scale electronic structure and transport calculations of organic polymers on the K computer," AIP Conf. Proc. in press; Preprint:http://arxiv.org/abs/1603.09616.

[7] A. Marek, V. Blum, R. Johanni, V. Havu, B. Lang, T. Auckenthaler, A. Heinecke, H.-J. Bungartz, and H. Lederer, "The ELPA Library – Scalable Parallel Eigenvalue Solutions for Electronic Structure Theory and Computational Science," J. Phys. Condens. Matter **26**, 213201, 2014; http://elpa.rzg.mpg.de/

[8] T. Imamura, Y. Hirota, T. Fukaya, S. Yamada, and M. Machida, "EigenExa: high performance dense eigensolver, present and future," *8th International Workshop on Parallel Matrix Algorithms and Applications (PMAA14), Lugano, Switzerland, 2014*; http://www.aics.riken.jp/labs/lpnctrt/index_e.html

[9] Y. Hasegawa, J.-I. Iwata, M. Tsuji, D. Takahashi, A. Oshiyama, K. Minami, T. Boku, F. Shoji, A. Uno, M. Kurokawa, H. Inoue, I. Miyoshi, and M. Yokokawa, "First principles calculation of electronic states of a silicon nanowire with 100,000 atoms on the K computer," *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis (SC11)*, Seattle, WA, 2011, Article No. 1.

[10] W. Kohn, "Density functional and density matrix method scaling linearly with the number of atoms," Phys. Rev. Lett. **76**, pp. 3168–3171, 1996.

[11] J. M. Soler, E. Artacho, J. D. Gale, A. García, J. Junquera, P. Ordejón, and D. Sänchez-Portal, "The SIESTA method for ab initio order-N materials simulation," J. Phys. Condens. Matter **14**, 2745, 2002.

[12] C.-K. Skylaris, P. D. Haynes, A. A. Mostofi, and M. C. Payne, "Introducing ONETEP: Linear-scaling density functional simulations on parallel computers," J. Chem. Phys. 122, 084119, 2005.

[13] T. Ozaki, "$\mathrm{O}(N)$ Krylov-subspace method for large-scale ab initio electronic structure calculations," Phys. Rev. B 74, 245101, 2006.

[14] M. J. Gillan, D. R. Bowler, A. S. Torralba, and T. Miyazaki, "Order-N first-principles calculations with the conquest code," Comp. Phys. Commun. **177**, 14, 2007.

[15] A. Frommer, "BiCGStab(l) for Families of Shifted Linear Systems," Computing **70**, pp. 87–109, 2003.

[16] R. Takayama, T. Hoshi, T. Sogabe, S.-L. Zhang, and T. Fujiwara, "Linear algebraic calculation of the Green's function for large-scale electronic structure theory," Phys. Rev. B 73, 165108, pp.1-9, 2006.

[17] T. Sogabe, T. Hoshi, S.-L. Zhang, and T. Fujiwara, "On a weighted quasi-residual minimization strategy of the QMR method for solving complex symmetric shifted linear systems," Electron. Trans. Numer. Anal. 31, pp. 126-140, 2008.

[18] H. Teng, T. Fujiwara, T. Hoshi, T. Sogabe, S.-L. Zhang, and S. Yamamoto, "Efficient and accurate linear algebraic methods for large-scale electronic structure calculations with nonorthogonal atomic orbitals," Phys. Rev. B 83, 165103, 12pp, 2011.

[19] T. Sogabe, T. Hoshi, S.-L. Zhang, and T. Fujiwara, "Solution of generalized shifted linear systems with complex symmetric matrices," J. Comp. Phys. 231, 5669-5684, 2012.

[20] S. Yamamoto, T. Sogabe, T. Hoshi, S.-L. Zhang, and T. Fujiwara, "Shifted Conjugate-Orthogonal-Conjugate-Gradient Method and Its Application to Double Orbital Extended Hubbard Model," J. Phys. Soc. Jpn., **77**, 114713, 2008.

[21] T. Mizusaki, K Kaneko, M. Honma, and T. Sakurai, "Filter diagonalization of shell-model calculations," Phys. Rev. C **82**, 024310, 2010.

[22] F. Giustino, M. L. Cohen, and S. G. Louie, "GW method with the self-consistent Sternheimer equation," Phys. Rev. E **81**, 115105, 2010.

[23] S. Iwase, T. Hoshi, and T. Ono, "Numerical solver for first-principles transport calculation based on real-space finite-difference method," Phys. Rev. E **91**, 063305, 2015.

[24] T. Adachi, N. Shida, K. Miura, S. Sumimoto, A. Uno, M, Kurokawa, F. Shoji, and M. Yokokawa, "The design of ultra scalable MPI collective communication on the K computer," Comput. Sci. Res. Dev. **28**, pp. 147–155, 2013.

[25] T. Yokota, P. Zalar, M. Kaltenbrunner, H. Jinno1, N. Matsuhisa, H. Kitanosako, Y. Tachibana, W. Yukita1, M. Koizumi1 and T. Someya, 'Ultra-flexible organic photonic skin', Sci. Adv. 2, e1501856, 2016.

[26] R. Dronskowski and P. E. Blöchl, "Crystal orbital Hamilton populations (COHP): energy-resolved visualization of chemical bonding in solids based on density-functional calculations," J. Phys. Chem. **33**, 97, 1993; `http://www.cohp.de/`